arXiv:1609.06840v1 [cs.LG] 22 Sep 2016

# Exact Sampling from Determinantal Point Processes

**Philipp Hennig**
Max Planck Institute for Intelligent Systems
Tübingen, Germany
ph@tue.mpg.de

**Roman Garnett**
Washington University in St. Louis
St. Louis, MO, USA
garnett@wustl.edu

## Abstract

Determinantal point processes (DPPs) are an important concept in random matrix theory and combinatorics. They have also recently attracted interest in the study of numerical methods for machine learning, as they offer an elegant "missing link" between independent Monte Carlo sampling and deterministic evaluation on regular grids, applicable to a general set of spaces. This is helpful whenever an algorithm *explores* to reduce uncertainty, such as in active learning, Bayesian optimization, reinforcement learning, and marginalization in graphical models. To draw samples from a DPP in practice, existing literature focuses on approximate schemes of low cost, or comparably inefficient exact algorithms like rejection sampling. We point out that, for many settings of relevance to machine learning, it is also possible to draw *exact* samples from DPPs on continuous domains. We start from an intuitive example on the real line, which is then generalized to multivariate real vector spaces. We also compare to previously studied approximations, showing that exact sampling, despite higher cost, can be preferable where precision is needed.

## 1 Introduction

Determinantal point processes (DPPs), introduced by Macchi [1975], are stochastic point processes whose joint probability measure is proportional to the determinant of a positive definite kernel Gram matrix (a more formal introduction follows in §2 below). Intuitively, this introduces a dependence between points sampled from such processes that gives them a "repulsive" property—point sets drawn from DPPs tend to cover a space more regularly than uniform random samples (cf. Figure 1).

DPPs initially arose in the study of fermionic gases in physics and have since found application in other areas, including random matrix theory [Mehta, 1991]. A review of their statistical properties is provided by Soshnikov [2000]. DPPs have seen less attention in machine learning and statistics than in physics, but DPPs on discrete domains have been used as diversity-inducing priors [e.g., Kulesza and Taskar, 2012]. An arguably under-explored direction is that they provide an elegant theoretical handle on the notion of *exploration* that is of interest across machine learning. In areas like active and reinforcement learning, as well as Bayesian optimization and numerical tasks like marginalization in graphical models, the basic challenge is that the algorithm should in some sense "probe" an input domain in a maximally informative way, while also ensuring that no area is ignored indefinitely.

As we will review in §2 below, DPPs have a direct connection to the variance function of Gaussian process regression models, which closely ties them to the basic probabilistic algorithms in the aforementioned fields. Of particular interest in the context of Bayesian optimization is the extensive work of Hough et al. [2009], which links the zero-crossings (e.g., the roots of gradients) of a particular class of Gaussian processes to DPPs. In the area of integration—in particular, marginalization in probabilistic models—there is already growing interest in DPPs. The two observations that independent (Monte Carlo) sampling leads to sub-optimal convergence rates of integration estimators [O'Hagan, 1987], and that classic quadrature rules can be explicitly interpreted as Gaussian process regression [Diaconis, 1988, O'Hagan, 1991], has sparked curiosity about structured kernel models

29th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

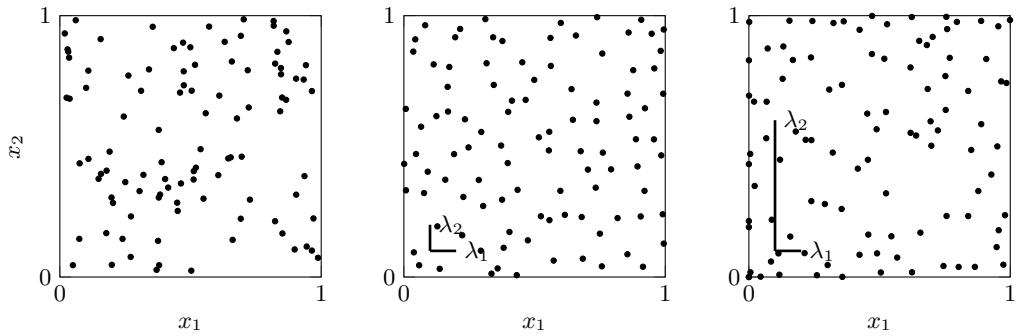


Figure 1: One hundred points sampled from, respectively, a uniform random (left) and two determinantal point processes (center, right) with kernel $k_{\text{SE}}$ (5) of different length scales. Uniform random numbers exhibit "clumping," randomly occurring patterns responsible for the stochastic convergence rate of Monte Carlo estimators. DPPs have a self-avoiding characteristic, yielding better coverage.

for Bayesian quadrature [Osborne et al., 2012b,a, Gunter et al., 2014, Briol et al., 2015]. DPPs have been suggested as exploration strategies for such models, an idea corroborated by a recent theoretical analysis by Bardenet and Hardy [2016], who show improved convergence rates for DPP integral estimators compared to Monte Carlo. We discuss this further in §2.1, where we also reinterpret an adaptive strategy for rapid Bayesian quadrature by Gunter et al. [2014] in terms of DPPs.

Despite these theoretical arguments in favor of DPPs in continuous spaces, this model class has found only limited practical use as a tool in machine learning and statistics, and almost entirely in discrete domains [e.g., Kulesza and Taskar, 2012, Kang, 2013]. One hurdle is that the algorithms typically proposed for sampling DPPs are either computationally taxing or approximate, in particular if the input domain is continuous and/or high-dimensional (more in §2.1 below). The goal of this paper is to point out that this is a historical legacy that should not hold back the use of DPPs as an algorithmic ingredient in machine learning: the historical uses of DPPs, primarily in physics, have model structure that makes exact and efficient sampling difficult. In machine learning, where the designer enjoys algebraic freedom to design the model, this is much less of a problem. On the other hand, numerical uses like the above often involve a comparably small sample set anyway for outside reasons, which makes inverting Gram matrices an acceptable cost. In §3, we present an analytic and exact sampling scheme for DPPs, by way of example on the most popular (Gaussian) kernel in machine learning (generalized in §3.4).

## 2 Determinantal point processes

Let $k\colon \mathbb{X} \times \mathbb{X} \to \mathbb{R}$ be a symmetric positive semi-definite kernel over some bounded space $\mathbb{X}$. We will use a notation familiar to readers of the kernel and Gaussian process literature. Given two sets $A := [a_1, \ldots, a_I], B := [b_1, \ldots, b_J] \subseteq \mathbb{X}$, the symbol $K_{AB} \in \mathbb{R}^{I \times J}$ is a matrix containing the elements $[K_{AB}]_{ij} = k(a_i, b_j)$. For our purposes, a determinantal point process is a stochastic process, such that a finite sample $X := [x_1, \ldots, x_N] \subset \mathbb{X}$ from the process has joint probability

$$p(x_1, \ldots, x_N) = \frac{Z}{N!} \det K_{XX}. \tag{1}$$

Here, $Z$ denotes a normalization constant, the value of which is of minor importance (its existence can be shown using a general argument [Hough et al., 2006]). More precise definitions can be found in Soshnikov [2000], Hough et al. [2009], and Bardenet and Hardy [2016]. They require a discussion of base measures and other properties of point processes, which unnecessarily complicate the exposition in our context. Kulesza and Taskar [2012] also provide a relatively complete introduction to discrete DPPs, where $\mathbb{X}$ is restricted to be a discrete space. We focus on the continuous case.

We make two observations in passing, which may be helpful for intuition. First, Eq. (1) clearly assigns zero measure to samples in which any two distinct samples have the same value, $x_i = x_j, j \neq i$. For such samples, there would be a permutation of $K_{XX}$ that contains at least one $2 \times 2$ minor containing exclusively $k(x_i, x_i)$, thus the determinant would vanish. Second, speaking somewhat

informally, uniform random point processes can be seen as a special case of DPPs for the Dirac kernel $k(x, x') = \delta(x - x')$. (There is a technical complication since even this kernel does not allow for pairwise identical samples, but such events also have measure zero under random sampling.)

### 2.1 Sampling from DPPs — the connection to active exploration

Equation (1) is an unwieldy object from an algorithmic perspective because it directly addresses the joint distribution of $N$ samples. To iteratively draw samples of increasing size, the conditional distribution $p(x_i \mid X_{1:i-1})$ is required, where $X_{1:i-1} := [x_1, \dots, x_{i-1}]$. By repeatedly applying the well-known determinant lemma $\det(Z + UWV^\intercal) = \det(Z)\det(W)\det(W^{-1} + V^\intercal Z^{-1} U)$, this conditional distribution can be found to be [Hough et al., 2006, Prop. 19]

$$p(x_1, \dots, x_N) = Z \prod_{i=1}^{N} p(x_i \mid X_{1:i-1}) = Z \prod_{i=1}^{N} \frac{1}{N - i + 1} \mathbb{V}_i(x_i), \tag{2}$$

using the function [see also Bardenet and Hardy, 2016]

$$\mathbb{V}_i(x) := \begin{cases} K_{xx} & \text{if } i = 1; \\ K_{xx} - K_{xX_{1:i-1}} K^{-1}_{X_{1:i-1}X_{1:i-1}} K_{X_{1:i-1}x} & \text{otherwise.} \end{cases} \tag{3}$$

This conditional function will be familiar to readers experienced with Gaussian process models: it is equal to the *posterior predictive variance* of a Gaussian process regression model conditioned on function values at $X_{1:i-1}$. In preparation for the derivations in §3, we introduce the shorthand $K_{(i)} := K_{1:i-1,1:i-1}$ and re-formulate Eq. (3) for $i > 1$ more explicitly as

$$\mathbb{V}_i(x) = k(x, x) - \sum_{a,b=1}^{i-1} k(x, x_a) k(x, x_b) [K^{-1}_{(i)}]_{ab}. \tag{4}$$

Because $\mathbb{V}_i(x)$ can be interpreted as the posterior variance of a Gaussian process model, one can think of a DPP as the point process arising from the following elementary *active learning* strategy: Consider an algorithm aiming to learn the function $f : \mathbb{X} \to \mathbb{R}$ by choosing evaluation points ("designs") $X$, using a Gaussian process prior $p(f) = \mathcal{GP}(\mu, k)$ with arbitrary mean function $\mu : \mathbb{X} \to \mathbb{R}$. Aiming to collect informative observations, the algorithm may adopt the policy to evaluate $f$ at a point $x_i$ with probability proportional to $\mathbb{V}_i(x_i)$. If it does so, its designs are samples from the DPP associated with $k$. This strategy can be motivated from within the Bayesian inference framework: evaluating a GP-distributed function at samples from the associated DPP amounts to drawing evaluation points with probability proportional to the expected information gain (evaluating at a mode of the DPP maximizes expected information gain about $f$). From outside the Bayesian framework, Bardenet and Hardy [2016] show that this policy is useful in so far as the resulting empirical estimator $\mathbb{E}_N[f]$ for expectations of $f$ (even if $f$ is *not* a true sample from $\mathcal{GP}(\mu, k)$, or even an element of the RKHS associated with $k$) converges at a rate dominating that of the Monte Carlo estimator.

In fact DPPs have already been used, albeit implicitly, in quadrature for machine learning: Gunter et al. [2014] proposed a fast method for Bayesian quadrature to estimate marginal likelihoods. To model the integrand as a strictly positive function, the square root of the likelihood was modeled with a GP. Given data, a second GP was fit to the likelihood, accounting for the nonlinear transformation via linearization or moment matching. This induced GP was used to estimate the desired integral via Bayesian quadrature, *and to choose evaluation points, by (approximately) maximizing the posterior variance of the transformed GP*. The WSABI-L method proposed in the op.cit. can be seen as choosing MAP samples from a DPP measure. Following Kulesza and Taskar [2012], the DPP kernel can be decomposed into a point-wise "quality" term and a normalized "diversity" kernel. Setting the former to the posterior mean of the underlying GP, and the latter to the posterior covariance, recovers the model of Gunter et al. [2014].

From a computational perspective, Eq. (2) poses two challenges, which seem to have tempered interest in sampling DPPs in continuous spaces, and numerical applications, so far:

1. For general kernels $k$, there is usually no analytic cumulative density function for $p(x_i \mid X_{1:i-1})$, which is required to draw exact samples from this distribution. Even elaborate studies of sampling methods from DPPs in continuous spaces [Scardicchio et al., 2009,

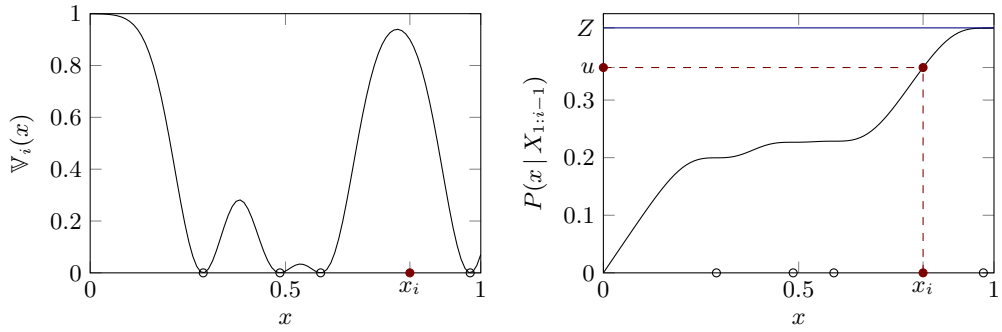


Figure 2: Sketch illustrating analytic sampling from a DPP in one dimension, using the square-exponential kernel (Eq. 5). Left: conditional probability $\mathbb{V}_i(x)$ for $i = 5$. Previous samples $X_{1:4}$ are drawn as empty circles; the new sample is shown in solid red. Right: The sample is drawn by computing the cumulative density $P$ (black line), drawing a scaled uniform random sample $u$ and finding the point $x_i$ such that $P(x_i) = u$, by interval bisection (Alg. 1, lines 9–12).

Bardenet and Hardy, 2016] thus rely on rejection sampling—at least for multidimensional domains—sometimes with carefully crafted proposal densities. This problem is pertinent in areas like physics where the kernel $k$ is predetermined by the problem of interest, and tends to have significant structure. This problem is much less severe in machine learning, because our community enjoys freedom in the design of models and can thus choose kernels with convenient analytic properties. Doing so directly yields an exact, efficient algorithm for the generation of samples from DPPs, even in high-dimensional domains $\mathbb{X}$.

2. Even if the kernel is analytically convenient, Eq. (3) involves the matrix inverse of $K_{(i)}$. Given the inverse of $K_{(i-1)}$ from the preceding step in the iterative sampling scheme, this inverse can be computed with complexity $\mathcal{O}\big((i-1)^2\big)$, using the matrix inversion lemma. Even so, the cost of drawing $N$ samples remains $\mathcal{O}(N^3)$. This issue is directly connected to inference in Gaussian process regression models, and many approximation schemes have been proposed in that area over the past decade. Affandi et al. [2013] proposed leveraging such fast approximation schemes to produce approximate DPP samples in $\mathcal{O}(N)$. We will show in §4 that this can introduce significant artifacts. In use cases like Bayesian optimization and quadrature, were the number $N$ of function evaluations is often low, and a Gram matrix is computed/inverted anyway, the cubic cost of exact sampling can be unproblematic, and precision may be more important.

## 3 Method

We point out an algorithm computing exact samples from a DPP if the kernel $k$ is analytically integrable. To ease intuition, the derivations will be by way of example, using the exceedingly popular square-exponential (aka. Gaussian, RBF) kernel $k_{\text{SE}} \colon \mathbb{R}^D \times \mathbb{R}^D \to \mathbb{R}$ over the real vector space[1]

$$k_{\text{SE}}(a, b) = \exp\left(-\frac{1}{2}\sum_{d=1}^{D}\frac{(a-b)_d^2}{\lambda_d^2}\right). \tag{5}$$

Assume that the sampling domain is the unit cube $x \in [0, 1]^D$ (for a more general box constraint $\tilde{x}_d \in [a_d, b_d]$ for each $d = 1, \ldots, D$, consider the linear transformation $x_d = {}^{\tilde{x}_d - a_d}/_{b_d - a_d}$). To simplify things even further, we initially consider the univariate problem, $D = 1$, then generalize to arbitrary dimensionality. The resulting algorithm draws $N$ samples at cost $\mathcal{O}(DN^2 + N^3)$. A general form of the algorithm is summarized in pseudo-code in Algorithm 1.

[1]There is a more-general version of this kernel using the Mahalanobis distance induced by a positive definite matrix $\Lambda \in \mathbb{R}^{D\times D}$. The algorithm as described below can be generalized to that form by rotating the input dimensions to the eigenvectors of $\Lambda$. The kernel $k$ can be scaled by an arbitrary positive scalar $\theta \in \mathbb{R}_+$, i.e., $k(x, x') \mapsto \theta k(x, x')$ without changing the DPP measure—this simply scales the normalization constant $Z$.

### 3.1 Sampling in one dimension

Exact samples can be drawn using the classic form of computing a non-normalized cumulative density function

$$P(x \,|\, X_{1:i-1}) = \int_0^x p(x \,|\, X_{1:i-1}) \,\mathrm{d}x, \tag{6}$$

and transforming standard uniform random variables $u \sim \mathbb{U}[0, P(1)]$, produced by a pseudo-random number generator, into exact samples from $p$, by setting (cf. Figure 2)

$$x = P^{-1}(u \,|\, X_{1:i-1}) = \{x \,|\, P(x) = u\}. \tag{7}$$

For the univariate square-exponential kernel (Eq. 5), Eq. (4) can be re-written, using standard properties of the Gaussian function, as

$$\mathbb{V}_i(x) = 1 - \sum_{a,b=1}^{i-1} \exp\left(-\frac{(x - m_{ab})^2}{\lambda^2}\right) \underbrace{\exp\left(-\frac{(x_a - x_b)^2}{4\lambda^2}\right)}_{=:M_{(i),ab}} [K_{(i)}^{-1}]_{ab}. \tag{8}$$

where $m_{ab} := {}^1\!/_2(x_a + x_b)$, and we have defined a matrix $M_{(i)} \in \mathbb{R}^{(i-1)\times(i-1)}$. Together, the variables $m, M, K^{-1}$ provide the "sufficient statistics" of the sample needed to draw the subsequent point. After $x_i$ has been drawn, these three variables can be updated in $\mathcal{O}(N^2)$—using the matrix inversion lemma to update $K_{(i+1)}^{-1}$; the other two variables can be updated in $\mathcal{O}(N)$. With these objects, the cumulative density is

$$P(x \,|\, X_{1:i-1}) = x - \frac{\sqrt{\pi}\lambda}{2} \sum_{a,b=1}^{i-1} \left[\operatorname{erf}\left(\frac{x - m_{ab}}{\lambda}\right) + \operatorname{erf}\left(\frac{m_{ab}}{\lambda}\right)\right] [M_{(i)} \odot K_{(i)}^{-1}]_{ab}. \tag{9}$$

Here, $\odot$ is the Hadamard (element-wise) product, and we have used $\operatorname{erf}(x) = -\operatorname{erf}(-x)$. Given a uniform random draw $u$, all that is left to do is to find $x$ such that $P(x \,|\, X_{1:i-1}) = u$. A straightforward, numerically robust, albeit not particularly ingenious way to do so is by interval bisection. A more elegant search strategy could be constructed using grid refinement methods similar to the popular Ziggurat algorithm of Marsaglia and Tsang [2000], which we skip here since our goal is merely to highlight generalizable structure, not to find an extremely efficient solution.

### 3.2 Multivariate samples

For square-exponential kernel DPPs in dimension $D > 1$, the function $\mathbb{V}_i(x)$ retains much of its structure. Equation (8) simply turns into (defining the elements of a new matrix $\mathbf{M} \in \mathbb{R}^{(i-1)\times(i-1)}$ analogous to $M$ in Eq. (8))

$$\mathbb{V}_i(x) = 1 - \sum_{a,b=1}^{i-1} \exp\left(-\sum_{d=1}^{D} \frac{(x - m_{ab})_d^2}{\lambda_d^2}\right) \underbrace{\exp\left(-\sum_{d=1}^{D} \frac{(x_a - x_b)_d^2}{4\lambda_d^2}\right)}_{=:\mathbf{M}_{(i),ab}} [K_{(i)}^{-1}]_{ab}. \tag{10}$$

The additional challenge in this multivariate case is to construct a parametrization of the cumulative density $P$. This step, too, can be performed in an iterative fashion, drawing one coordinate of the sample point $x_i$ after another (cf. Figure 3). Given that the first $d-1$ elements of $x_i$ are given by $x_{i,1:d-1}$, the cumulative density associated with the $d$th dimension is given by the sum rule:

$$P(x_{i,d} \,|\, X_{1:i-1}, x_{i,1:d-1}) = \int_0^{x_{i,d}} \int \cdots \int_0^1 p\big([x_{i,1:d-1}, \tilde{x}_{i,d}, \tilde{x}_{i,d+1:} \,|\, X_{1:i-1}\big) \,\mathrm{d}\tilde{x}_{i,d} \prod_{\tilde{d}=d+1}^{D} \mathrm{d}\tilde{x}_{i,\tilde{d}}. \tag{11}$$

For the square-exponential kernel, this works out to

$$\begin{aligned} P(x_{i,d} \,|\, X_{1:i-1}, x_{i,1:d-1}) = x - \sum_{a,b=1}^{i-1} \Bigg\{ &\exp\left(-\sum_{r=1}^{d-1} \frac{(x - m_{ab})_r^2}{\lambda_r^2}\right) [\mathbf{M}_{(i)} \odot K_{(i)}^{-1}]_{ab} \\ &\cdot \left(\prod_{\ell=d}^{D} \left(\operatorname{erf}\left(\frac{[x - m_{ab}]_\ell}{\lambda_\ell}\right) + \operatorname{erf}\left(\frac{[m_{ab}]_\ell}{\lambda_\ell}\right)\right) \frac{\sqrt{\pi}\lambda_\ell}{2}\right) \Bigg\}. \end{aligned} \tag{12}$$

Algorithm 1 provides a pseudo-code summary.

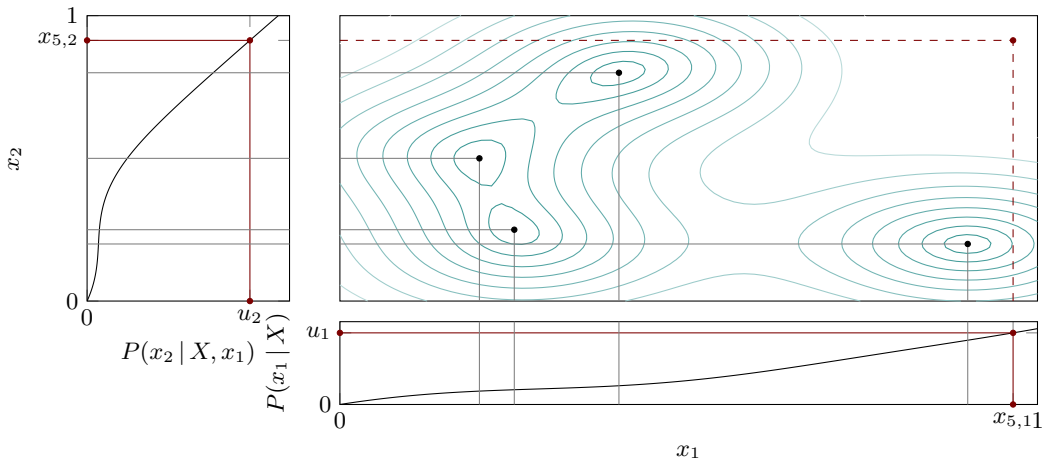


Figure 3: Drawing samples from a DPP in two dimensions (an example sufficiently general to extend to arbitrary dimensionality). Top right: surface plot of the multivariate probability density $\mathbb{V}_5(x)$. The preceding four samples are shown as black points. Bottom: The first coordinate of the fifth sample is drawn first, by computing the marginal density along this dimension. Left: The second element is then drawn by computing a cumulative density *conditioned* on the value of the first element.

**Algorithm 1** Exact sampling from DPPs with analytic kernels, on $[0, 1]^D$.

**procedure** DRAWFROMDPP($k, D, N$)
  $X \leftarrow \varnothing, m \leftarrow \varnothing, M \leftarrow \varnothing, K^{-1} \leftarrow \varnothing$ // initialize statistics of sample as empty
  **for** $n = 1, \ldots, N$ **do** // draw samples iteratively
    $x_n \leftarrow \varnothing$ // initialize current sample point
    **for** $d = 1, \ldots, D$ **do** // draw dimensions iteratively
      $P \leftarrow$PCONSTRUCT($m, M, K^{-1}$) // construct function for Eq. (12)
      $u \leftarrow P(1)\cdot$RAND($\cdot$) // draw scaled unit random number
      $I \leftarrow [0, 1]$ // initialize search interval
      **while** $|I| > \varepsilon$ **do** // bisection search
        $\mu \leftarrow 1/2(I_0 + I_1)$ // interval midpoint
        $I \leftarrow$ **if** $(P(\mu) < u)$ **then** $[\mu, I_1]$ **else** $[I_0, \mu]$ // bisect
      **end while**
      $x_n \leftarrow [x_n, \mu]$ // store sampled element
    **end for**
    $(X, m, M, K^{-1}) \leftarrow$UPDATESTATS($x_n, X, m, M, K^{-1}$) // update sample statistics (rank-1)
  **end for**
**end procedure**

### 3.3 Relation to previous work

The main goal of this paper is to point out that applications of DPPs in machine learning and numerics, thanks to their freedom to choose analytically convenient kernels, can utilize exact, computationally efficient sampling schemes in virtually arbitrary dimensionality. The algorithm presented here can be seen as a concrete realization of an abstract recipe introduced by Hough et al. [2006, Prop. 19]. Other literature on sampling discrete DPPs can be traced back to this work, including [Kulesza and Taskar, 2012, Kang, 2013, Wang et al., 2014]. This recipe was investigated numerically by Scardicchio et al. [2009]. However, these authors, restricted by the algebraic structure of their physical application, could only draw analytic samples in 1D and had to resort to rejection sampling in the multivariate case. Although rejection sampling with a decent proposal distribution can scale up to several dimensions, its computational cost rises exponentially with dimensionality. For the high-dimensional domains typical of machine learning problems, only exact sampling is practical.

The work most-related to ours is by Affandi et al. [2013], who considered *approximate* sampling of DPPs in continuous domains. Here we provide an exact algorithm that should be of interest for Bayesian optimization and integration. We explicitly compare between the quality of exact and approximate samples on a continuous domain §4 below.

### 3.4 Other analytical kernels

While the Gaussian kernel is the most widely used kernel in machine learning, it has some shortcomings, primarily that it makes very strong smoothness assumptions that can lead to instability in interpolation models. But with some algebraic elbow grease, the scheme of Eq. (11) can be extended to many other popular kernels, assuming they factorize,

$$k(a,b) = \prod_d^D k(a_d, b_d), \tag{13}$$

and the indefinite integrals

$$\int k(a,a)\,\mathrm{d}a \qquad \text{and} \qquad \int k(a,b)k(a,c)\,\mathrm{d}a \tag{14}$$

are analytically solvable. For example, the above results are applicable to the Matérn class of kernels [Stein, 1999] (including the exponential kernel, which induces the Ornstein-Uhlenbeck process), noting that, assuming w.l.o.g. $x_0 < a < b < x_1$,

$$\int_{x_0}^{x_1} \exp\left(-|x-a|\right)\exp\left(-|x-b|\right)\,\mathrm{d}x = \frac{e^{-a-b}}{2}\left(e^{2a} - e^{2x_0}\right) + (b-a)e^{a-b} + \frac{e^{a+b}}{2}\left(e^{2x_1} - e^{2b}\right), \tag{15}$$

and using results such as [see e.g., Gradshteyn and Ryzhik, 2007, §2.322]

$$\int xe^{ax}\,\mathrm{d}x = e^{ax}\left(\frac{x}{a} - \frac{1}{a^2}\right), \quad \int x^2e^{ax}\,\mathrm{d}x = e^{ax}\left(\frac{x^2}{a} - \frac{2x}{a^2} + \frac{2}{a^3}\right), \ldots \tag{16}$$

## 4 Comparison to finite-rank approximations

Due to the matrix inverse $K_{(i)}^{-1}$ in Eq. (4), the cost of evaluating $\mathbb{V}_N(x)$ grows cubically, $\mathcal{O}(N^3)$, with the sample size $N$. If large samples are required, an approximate approach of $\mathcal{O}(N)$ may be more appealing. In 2013, Affandi et al. proposed the use of standard low-rank approximations of the kernel for this purpose. We briefly review this idea here, and compare it empirically to the exact sampler below. If the kernel can be approximated well by a finite-rank expansion

$$k(a,b) \approx \sigma^2\boldsymbol{I} + \sum_{f,g=1}^{F} \phi_f(a)\Sigma_{fg}\phi_g(b), \tag{17}$$

using a collection of (not necessarily orthogonal) feature functions $\phi\colon \mathbb{X} \to \mathbb{R}$ and a symmetric positive definite matrix $\Sigma \in \mathbb{R}^{F\times F}$, then $\mathbb{V}$ can be approximated in $\mathcal{O}(F^3 + NF^2)$ time (i.e., linear in the sample size) as

$$\mathbb{V}(x) = \sum_{f,g=1}^{F} \phi_f(x)\phi_g(x)\left[\Sigma^{-1} + \sigma^{-1}\Phi\Phi^\intercal\right]_{fg}, \tag{18}$$

using the matrix $\Phi \in \mathbb{R}^{F\times N}$ with elements $\Phi_{fi} = \phi_f(x_i)$. This includes the case of degenerate kernels, i.e., where $k$ is *exactly* captured by such a finite-rank expansion, such as in simple linear and polynomial regression. Approximate cases include the Nyström approximation [Williams and Seeger, 2001], and spectral expansions [Rahimi and Recht, 2008]. The latter approach works for any translation invariant kernel over the unit hypercube (including the square-exponential), since the necessary eigenfunctions are the trigonometric functions. The necessary integrals in Eq. (14) can be solved using identities like [cf. Gradshteyn and Ryzhik, 2007, §2.532, for a more complete list]

$$\int \cos(ax)\cos(bx)\,\mathrm{d}x \overset{a\neq b}{=} +\frac{\sin((a-b)x)}{2(a-b)} + \frac{\sin((a+b)x)}{2(a+b)} \quad \text{or} \overset{a=b}{=} \frac{x}{2} + \frac{\sin(2ax)}{4a}. \tag{19}$$

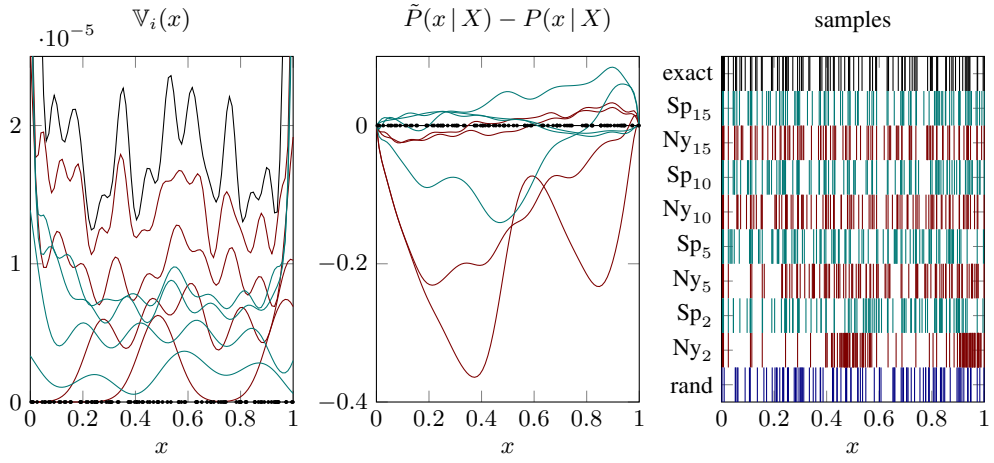


Figure 4: Comparison between exact and approximate sampling schemes. Left: probability density $\mathbb{V}_i(x)$ after $i = 100$ exact samples (black dots on abscissa). Exact density in black on top, approximations for Nyström (red) and spectral (green) approximations with $F = 5, 10, 15$ (ordering is obvious, quality increases with $F$). Center: Deviation of the associated approximate cumulative densities $\tilde{P}$ from the exact $P$. For this plot, all $\tilde{P}$ and $P$ were normalized (i.e. $\tilde{P}(1) = 1$) for easier comparison, and to clarify that the absolute value of $\mathbb{V}$ is less important than its shape. Right: Sample locations drawn exactly (black, top), and from the approximations. The different samplers used the same random seeds $u$ (shown for comparison in blue at the bottom), so if the approximations were exact, the black and red lines would match. All methods share the first 20 (exact) samples.

Figure 4 empirically compares samples drawn from a univariate DPP using these two approximations with exact samples drawn using the algorithm outlined above. It shows both the approximations to the density $\mathbb{V}$, the differences of the induced normalized cumulative density functions to the exact one, and approximate DPP samples, for both the Nyström and spectral approximations of varying fidelity $F$ (further details in caption). Rough approximations (i.e. with small $F$) can perform badly—in fact, the samples drawn this way may be *more* clumped than uniform random samples. This experiment was deliberately performed in 1D to take an optimistic stance on the potential quality of the approximations: For DPP sampling, the number of features necessary to cover higher-dimensional spaces increases exponentially with dimension. (Empirical evaluations of low-rank approximations for kernel regression show that they can work well even in multivariate settings [e.g., Rahimi and Recht, 2008], but empirical datasets tend to lie on low-dimensional manifolds, while DPP samples, by construction, cover the entire domain, thus require the approximation to actually cover the space.)

There is thus a nontrivial trade-off between cost and precision when sampling from a DPP. If a moderate number of samples (say, $N \lesssim 10^3$) are to be drawn in a space of high dimension $D$, it may be a better idea to draw exact samples at cost $\mathcal{O}(N^3 + N^2 D)$, rather than approximate samples to rank $F$ at cost $\mathcal{O}(NF^2 + F^3 + F^2 D)$, if $F$ scales exponentially with $D$. This is particularly true if, as in Bayesian optimization and quadrature, the reason to draw the samples is to condition a GP regression model that already requires the very same cubic cost matrix inversion anyway.

## 5 Conclusion

We have pointed out that, for popular kernel classes of machine learning and statistics, it is often possible to analytically draw exact samples from determinantal point processes on continuous spaces at comparably low computational cost. More precisely, this is true if the kernel defining the DPP has certain analytic properties defined in Eqs. (13) & (14). If the resulting sample locations are used to condition a kernel regressor, the sampling algorithm's cost is dominated by a matrix inversion that is required by the regressor anyway, with a very small remaining cost overhead. This immediate association between a kernel regressors and a DPP, coupled with the interesting theoretic properties of DPPs, make these point processes an interesting framework for areas like Bayesian optimization and quadrature. The algorithm outlined above provides the necessary tool to utilize this opportunity.

## References


R. H. Affandi, E. Fox, and B. Taskar. Approximate inference in continuous determinantal processes. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26, pages 1430–1438. 2013.

R. Bardenet and A. Hardy. Monte Carlo with determinantal point processes. *ArXiv e-print*, 1605.00361, May 2016.

F.-X. Briol, C. Oates, M. Girolami, and M. A. Osborne. Frank-Wolfe Bayesian quadrature: Probabilistic integration with theoretical guarantees. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 1162–1170. 2015.

P. Diaconis. Bayesian numerical analysis. *Statistical decision theory and related topics*, IV(1):163–175, 1988.

I. Gradshteyn and I. Ryzhik. *Table of Integrals, Series, and Products*. Academic Press, 7th edition, 2007.

T. Gunter, M. A. Osborne, R. Garnett, P. Hennig, and S. J. Roberts. Sampling for inference in probabilistic models with fast Bayesian quadrature. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, pages 2789–2797. 2014.

J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág. Determinantal processes and independence. *Probability Surveys*, 3:206–229, 2006.

J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág. *Zeros of Gaussian analytic functions and determinantal point processes*, volume 51, University Lecture Series. American Mathematical Society Providence, RI, 2009.

B. Kang. Fast determinantal point process sampling with application to clustering. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 26, pages 2319–2327. 2013.

A. Kulesza and B. Taskar. Determinantal point processes for machine learning. *Foundations and Trends in Machine Learning*, 5:123–286, 2012.

O. Macchi. The coincidence approach to stochastic point processes. *Advances in Applied Probability*, pages 83–122, 1975.

G. Marsaglia and W. W. Tsang. The Ziggurat method for generating random variables. *Journal of Statistical Software*, 5(8):1–7, 2000.

M. L. Mehta. *Random Matrices*. Academic Press, 1991.

A. O'Hagan. Monte Carlo is fundamentally unsound. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 36(2/3):247–249, 1987.

A. O'Hagan. Bayes–Hermite quadrature. *Journal of Statistical Planning and Inference*, 29(3):245–260, 1991.

M. Osborne, R. Garnett, Z. Ghahramani, D. K. Duvenaud, S. J. Roberts, and C. E. Rasmussen. Active learning of model evidence using Bayesian quadrature. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 46–54. 2012a.

M. Osborne, R. Garnett, S. Roberts, C. Hart, S. Aigrain, and N. Gibson. Bayesian quadrature for ratios. In M. Girolami and N. Lawrence, editors, *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 15, pages 832–840, 2012b.

A. Rahimi and B. Recht. Random features for large-scale kernel machines. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20, pages 1177–1184. 2008.

A. Scardicchio, C. E. Zachary, and S. Torquato. Statistical properties of determinantal point processes in high-dimensional Euclidean spaces. *Physical Review E*, 79(4):041108, 2009.

A. Soshnikov. Determinantal random point fields. *Russian Mathematical Surveys*, 55(5):923–975, 2000.

M. L. Stein. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer S&B M, 1999.

S. Wang, C. Zhang, H. Qian, and Z. Zhang. Using the matrix ridge approximation to speedup determinantal point processes sampling algorithms. In C. E. Brodley and P. Stone, editors, *AAAI Conference on Artificial Intelligence*, volume 28, pages 2121–2127. 2014.

C. K. I. Williams and M. Seeger. Using the Nyström method to speed up kernel machines. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems*, volume 13, pages 682–688. 2001.